

Dashboard & Reporting

Using the Cyber Crucible web application - agent visibility, process activity, memory analytics, and executive reporting.

- [Can Cyber Crucible silently respond to an attack, but not alert/notify me?](#)
- [How can I investigate the root cause of an event?](#)
- [What does the Process Injections Page show?](#)
- [What does the Process Creations Page show?](#)
- [What are Executive Report Summaries?](#)
- [How do I show hidden agents?](#)
- [How do I hide or remove an agent?](#)
- [If I hide an agent, will I still get alerts & notifications?](#)

Can Cyber Crucible silently respond to an attack, but not alert/notify me?

You may find that there are legitimate scenarios where you do not want an application doing something (probably on your machines, but do not need notified every time).

This is a separate situation as to where you need something whitelisted, where no action is taken by Cyber Crucible's data extortion prevention.

The capability you are looking for is silent automated responses.

Making an existing automated response silent in the future:

First, go to the Manage Incidents page. We're changing the name to Manage Responses in the near future.

Second, click on the speaker/volume button as seen below.

4784131.png?width=680

Clicking this button brings up the "Silence Incident" or "Silence Auto-Response" modal.

We have analytics that check whether a response should be silent based on previously selected programs and arguments.

In this case, we see that Adobe runs a web server as part of its software suite, and we do not want any future alerts that the web server was suspended.

(In this case, we inserted malicious javascript into the file that Adobe ran, and Cyber Crucible suspended Adobe's web server from stealing or encrypting data.)

4784138.png

Click on the "Take me to Silent Responses Page"

There will be a menu up for inserting that specific silence rule, but let's look at the page first:

5013506.png?width=680

We see here that the rules are based on both the process and the process arguments for an automated response, that should not be included in notification emails or other alerts.

The suspending still happens - just no notification about it.

We see here that the we do not want the Microsoft Edge browser, or the Chrome browser, to be able to run when using those arguments, and wildcards before and after to account for other arguments that are present.

(Security explainer insight: Both browsers run an identical piece of code, that loads the browsers in the background without a window on the screen, in order to scan and read your files behind the scenes. We do not want our web browsers doing that behind our backs!)

Regardless of whether you are creating an entirely pre-emptive silent response rule or not, the same window below (in the next section) is the next step.

Creating a Silent Response

Here we see both the program path, and the program arguments. Wildcards are allowed, to ensure something like a username or application version number won't stop your behavioral rule from working.

An example of why wildcards are needed is because, there might be a program on all of your employee desktops that are in a group.

C:\Users\mary\Desktop\runme.exe

C:\Users\joe\Desktop\runme.exe

The solution would be:

C:\Users*\Desktop\runme.exe

5046277.png

Here, we see that we are creating a silent response for the entire TEMPORARY GIRAFFE 2 group, for an Adobe Created Cloud path, and we only want to silence responses that have a specific argument (in this case a file path to a certain javascript file).

There are no unique pieces to the path, like usernames, so we can create the rule as is.

Responses may be named, with the names being unique per group.

The rules start being distributed to agents within minutes.

How can I investigate the root cause of an event?

- [Memory Analytics](#)
- [Process Injection Analytics](#)
- [Process Creation Analytics](#)
- [Putting it altogether - an end to end investigation](#)
 - [Why did we choose to show this example?](#)

There are three building blocks of analysis available to you, while investigating an alert.

It is currently more likely that an attacker will conduct a variety of activities on a system, using a combination of in-memory techniques and process hijacking (injection or hollowing), than running an obvious “hacker.exe” program.

Memory Analytics

The first building block is found on the Response or Incident page itself. Memory analytics track the behaviors of the process in memory during continuous inspection of possible data extortion (theft or encryption) behaviors.

A memory state that indicates likely tampering is typically indicative of an issue.

It isn't always malicious on itself, as poor programming practices or software bugs can also result in the memory state of a running program to be tampered with.

63504385.png?width=680

Here we see an instance where an Adobe Acrobat process being exhibit possible data extortion behaviors after opening a downloaded PDF, and we see the Acrobat process memory had been modified. The Acrobat.exe process on the disk has not been affected in any way.

Forensics analysis of the memory involved with the suspended Acrobat.exe will likely indicate an issue of concern - possibly an exploit in use, if not process injection or process hollowing.

For extortion responses with a Modified Memory value of True, there is a download icon in the Root Cause Analysis column to download the Memory Diff file for the response:

63602693.png?width=680

The Memory Diff File is the compiled source code of what was injected into the program. It likely contains malware and/or exploit code, for a malware or security analyst to examine. Note that this file does not contain the entire program, just the part that was altered through an exploit or attach technique like process injection. Learn more about this file [here](#).

Process Injection Analytics

The first indicator that there may be a process injection event relevant to an automated response is the Untrusted Remote Threads column for an Incident/Response. This indicates that kernel-level behavioral analysis revealed an untrusted remote thread being observed with the process conducting data theft or ransomware-like behavior.

63471635.png?width=680

Further investigation may be performed on the Process Injection page. Untrusted Remote Threads reading True is a good indicator whether to go here.

The Root Cause Analysis column also has an icon to automatically redirect you to the Injections Page to only show the related process injections to the response

63602703.png?width=680

After clicking the injections icon, you will be taken to the Process Injections page where you will see only the related process injections to the response:

63569930.png?width=680

Here in the Process Injection Events page, we see the injector process path and arguments, and the injectee path and arguments. Usually, the arguments play a very important role in memory analytics, or indicate, for instance, what powershell commands to run. The process ID's are present as well, to assist in tracing iterative injections done as an attacker hops from process to process, or for looking up root cause analysis in the Process Creation page.

63438879.png?width=680

How did we know rundll32.exe was really an antivirus vendor? Read on!

Process Creation Analytics

It is very rare that an attacker only uses one process during their attack. They may not even know they tools they are using are spawning new processes, and using various administrative tools on the infected system.

The Process Creation page tracks all processes that open other processes (or the same), the arguments used, and the process ID's, or Pids, for each source and destination.

This is performed at the kernel level, meaning nothing is missed.

Note: Microsoft telemetry has some of this data, but not all. We elect to not use that data source, because we've observed that telemetry source has been seen to miss important events, or be switched off by attackers during an attack.

How good is our visibility?

Putting it altogether - an end to end investigation

How good is our visibility, and how fast can you perform Root Cause Analysis? Check this example out - an antivirus with arguably some of the highest permissions on a system, was observed trying to inject code into our Cyber Crucible software. We hope it was for investigating our software, but we still go into self-protection mode (so, this caused our software to reject the antivirus' tampering...there are no back doors allowing access to edit or use our software by external parties).

So, let's start here, with an automated response:

63569937.png?width=680

Obviously Cyber Crucible is not ransomware. We see here that a process injection attempt involving suspicious remote thread behavior, in conjunction with data extortion-like behaviors.

OK - let's look at what's going on in Process Injection Events by clicking the injections icon . 5 seconds later - we have our answer!

63471644.png?width=56463438879.png?width=680

But wait - rundll? That's not a hacker, is it? Surely, that's not a hacker.exe program. Let's do some filtering, to get our answer on the Process Creation page. We don't have a screenshot here, but process ID's are available for all responses and process injections. Let's go back to the responses page, and click the process creations icon to take us to the Process Creations page and automatically filter the grid to show the related process creations of this response.

63537184.png63438886.png?width=680

Quick and easy, we now see that the "root" (no pun intended) was WebRoot, an antivirus vendor, spawned a rundll32.exe process to load one of WebRoot's DLL's, which then tried to do *something* with Cyber Crucible's software, to force it to iterate through files like a data theft tool.

Why did we choose to show this example?

We have copious examples of attacker tradecraft, with varying levels of competence on the attackers' part. Antiviruses achieve a level of permissions on a system that require a very high level of skill, technology, and preparation by attackers. Cyber Crucible detecting “hacker like” behaviors from a mainstream security tool, then automatically stopping it, demonstrates some of the highest levels of engineering and security excellence on our part, without getting lost in individual exploits or attacker tradecraft.

What does the Process Injections Page show?

Process injections are a software engineering capability that can appear in non-malicious context, or directly used by an attacker. In fact, process injection and associated techniques are a favorite of attackers, for a variety of reasons outside the scope of this article (but if you ask, we can point you towards training).

Cyber Crucible exposes process injection events at the kernel level. That means you have complete visibility for all processes and their information, regardless of the permissions of that process. Something we've found is that once attackers or software reaches a certain level of permissions, a lot of the traditional sources of telemetry for security tools go silent. Never missing an important piece of data, and never going silent, is a big part of our zero trust product design.

On the Process Injections page, you will see the following per event, with filters available to refine your queries:

[CC Showing All Process Injection Columns.mp4](#)

Normally, in the course of an investigation, you'll find that there is additional activity in the Process Creation page.

What does the Process Creations Page show?

It is very rare, possibly never, that an attacker performs an entire attack from within one process.

The fact that an attack tool or a Windows library is using multiple processes and Windows programs is sometimes invisible to users of the malware (aka, the criminals).

Process behaviors represent an important source of variables during Cyber Crucible decision making, which involves behavioral indicators from a variety of sources.

Sometimes these process behaviors may ultimately result in a process suspension from data theft or ransomware encryption, and can be traced from attacker entry point through the system, to ultimate extortion software execution.

Other times, process tracing can identify unauthorized remote access, fraud, insider threat, or unexpected software, but not data extortion. So, not something that Cyber Crucible's data extortion prevention software will automatically suspend, but still something you want to know about.

Whether the process data is part of a data extortion attack, something else of note, or just program behavior, the programs source and destination information, to include program arguments, are made available. This provides a rich set of data to trace attackers, including ones leveraging "[living off the land](#)" tactics.

Below is a screenshot of the type of data that can be gathered from process creations.

[CC_process_creation_demo.mp4](#)

What are Executive Report Summaries?

- [Introduction](#)
- [How to Opt-in and Opt-out of Receiving Reports](#)
- [Sample Report Content](#)

Introduction

Executive Report Summaries are reports emailed out to users that detail information about users' licenses, agents, and machine behaviors. The reports contain a summary on all groups the user is in, followed by a summary for each singular group.

Users have the option to receive weekly and monthly reports, as well as choose the format for the report sent in the email (HTML, ZIP, etc.).

How to Opt-in and Opt-out of Receiving Reports

After logging into our website and navigating to the Account Settings page, users will see the options to set their preferences for receiving weekly and monthly reports, as well as setting the report format that is sent in the email.

24051723.png?width=442

Sample Report Content

23887914.png?width=44224084489.png?width=44223855156.png?width=442

How do I show hidden agents?

Travel to the Agents page and locate the Dashboard Visibility column.

86048770.png?width=680

The default filter for this column is for hidden agents to not be visible.

To show Hidden agents, open the filter for this column and check the box for Hidden, and uncheck the box for Visible, then click Apply.

This will show only the Hidden agents. From there, you may conduct normal agent management operations, including changing visibility.

86016006.png?width=45386016013.png?width=680

How do I hide or remove an agent?

It is important to understand that removing an agent would also delete all of the data collected via that agent. That information may be useful someday for a future investigation.

Instead, hiding the agent from view, notifications, or alerts is probably the desired action.

There are two ways to do this.

Bulk Update Agents

The first method is easiest if there are multiple agents that need to be hidden from the web app.

First, select the agents you desire to be hidden, from the Agents page.

Normally, users will use filters based on installation status, machine name, or IP address first.

Second, select the “Bulk Set Agent Default Visibility” icon.

86048779.png?width=680

A modal will appear where you can select whether the selected agents' default visibility will be hidden or visible. Then, click the “Update Selected Agents” button.

86081547.png?width=448

The default setting for security notifications is to send offline agent alerts for hidden agents. See more about security notifications [here](#).

Update a Single Agent

On the Agents page, locate the Dashboard Visibility Column and toggle the Visible slider for the desired agent to change the visibility.

86016024.png?width=680

If I hide an agent, will I still get alerts & notifications?

The default setting for notifications is to send offline agent alerts for hidden agents.

We see below the default setting during notification creation:

Create New Security Notification

Group

Select

☒

 Send Offline Agent Alerts for Hidden Agents

☐

 Do Not Send Ransomware Activity Alerts for Silent Responses

☒

 Notify All Group Members

Name for Notification

Optional

Notification type:

☐ Ransomware Activity

☐ Offline Agent

☐ Abnormal Identity Access

Time Interval

Select

Create

Users normally make a separate notification category for hidden agent notifications, for the rare instances where that is desired. We recommend these notifications have different destination email aliases and possibly interval of notification.

If an existing notification requires an update to this field, travel to the Security Notifications page.

Scroll right on the grid until you see the Send Offline Alerts For Hidden Agents column, then toggle the slider as desired.

The change will take effect immediately.